

Monitor and Troubleshoot ASA Performance Issues

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Troubleshoot Performance Issues](#)

[Speed and Duplex Settings](#)

[CPU Utilization](#)

[High Memory Utilization](#)

[PortFast, Channeling, and Trunking](#)

[Network Address Translation \(NAT\)](#)

[Syslogs](#)

[SNMP](#)

[Reverse DNS Lookups](#)

[Show Commands](#)

[Show CPU Usage](#)

[Show Traffic](#)

[Show Perfmon](#)

[Show Blocks](#)

[Show Memory](#)

[Show Xlate](#)

[Show Conn Count](#)

[Show Interface](#)

[Show Processes](#)

[Command Summary](#)

[Related Information](#)

Introduction

This document describes the commands to use to monitor and troubleshoot the performance of a Cisco Adaptive Security Appliance (ASA).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on a Cisco Adaptive Security Appliance (ASA) that runs version 8.3 and later.


The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

Troubleshoot Performance Issues

In order to troubleshoot performance issues, check the basic areas described in this section.

 **Note:** If you have the output of the `show` command from your Cisco device, you can use the [Cisco CLI Analyzer](#) in order to display potential issues and fixes. The Cisco CLI Analyzer supports certain `show` commands. If you use the Cisco CLI Analyzer, you must be a registered Cisco user, you must be logged in to your Cisco account, and you must have JavaScript enabled within your browser.

Speed and Duplex Settings

The security appliance is preconfigured to autodetect the speed and duplex settings on an interface. However, several situations exist that can cause the automatic negotiation process to fail, which results in either speed or duplex mismatches (and performance issues). For mission-critical network infrastructure, Cisco manually hardcodes the speed and duplex on each interface so there is no chance for error. These devices generally do not move around, so if you configure them properly, you do not need to change them.

On any network device, link speed can be sensed, but duplex must be negotiated. If two network devices are configured to automatically negotiate speed and duplex, they exchange frames (called Fast Link Pulses, or FLPs) that advertise their speed and duplex capabilities. In order to a link partner that is not aware, these pulses are similar to regular 10 Mbps frames. In order to a link partner that can decode the pulses, the FLPs contain all the speed and duplex settings that the link partner can provide. The station that receives the FLPs acknowledges the frames, and the devices mutually agree on the highest speed and duplex settings that each can achieve. If one device does not support automatic negotiation, the other device receives the FLPs and transitions to parallel detection mode. In order to sense the speed of the partner, the device listens to the length of pulses, and then sets the speed based on the length. The problem arises with the duplex setting. Because duplex must be negotiated, the device that is set to automatically negotiate cannot determine the settings on the other device, so it defaults to half-duplex, as stated in the IEEE 802.3u standard.

For example, if you configure the ASA interface for automatic negotiation and connect it to a switch that is hardcoded for 100 Mbps and full-duplex, the ASA sends out FLPs. However, the switch does not respond because it is hardcoded for speed and duplex and does not participate in automatic negotiation. Because it receives no response from the switch, the ASA transitions into parallel detection mode and senses the length of the pulses in the frames that the switch sends out. That is, the ASA senses that the switch is set to 100 Mbps, so it sets the interface speed based on this. However, because the switch does not exchange FLPs, the ASA cannot detect if the switch can run full-duplex, so the ASA sets the interface duplex to half-duplex, as stated in the IEEE 803.2u standard. Because the switch is hardcoded to 100 Mbps and full-duplex, and the ASA has just automatically negotiated to 100 Mbps and half-duplex (as it does), the result is a duplex mismatch that can cause severe performance problems.

A speed or duplex mismatch is most frequently revealed when error counters on the interfaces in question increase. The most common errors are frame, cyclic redundancy checks (CRCs), and runts. If these values increment on your interface, either a speed/duplex mismatch or a cabling issue occurs. You must resolve this issue before you continue.

Example

<#root>

```
Interface GigabitEthernet0/0 "outside", is up, line protocol is up
  Hardware is i82546GB rev03, BW 1000 Mbps, DLY 10 usec
    Auto-Duplex(Full-duplex), Auto-Speed(100 Mbps)
    Input flow control is unsupported, output flow control is unsupported
    MAC address 0013.c480.b2b8, MTU 1500
    IP address 192.168.17.4, subnet mask 255.255.255.0
    311981 packets input, 20497296 bytes, 0 no buffer
    Received 311981 broadcasts,
```

```
157 runts
, 0 giants

379 input errors, 107 CRC, 273 frame
, 0 overrun, 0 ignored, 0 abort
  0 pause input, 0 resume input
  0 L2 decode drops
  121 packets output, 7744 bytes, 0 underruns
  0 pause output, 0 resume output
  0 output errors, 0 collisions, 1 interface resets
  0 late collisions, 0 deferred
  0 input reset drops, 0 output reset drops, 0 tx hangs
  input queue (blocks free curr/low): hardware (255/249)
  output queue (blocks free curr/low): hardware (255/254)
```

CPU Utilization

If you noticed the CPU utilization is high, complete these steps in order to troubleshoot:

1. Verify that the connection count in `show xlate count` is low.
2. Verify that the memory block is normal.
3. Verify that the number of ACLs is higher.
4. Issue the `show memory detail` command and verify that the memory used by the ASA is normal utilization.
5. Verify that the counts in `show processes cpu-hog` and `show processes memory` are normal.
6. Any host present inside or outside the security appliance can generate the malicious or mass traffic that can be a broadcast/multicast traffic and cause the high CPU utilization. In order to resolve this issue, configure an access list to deny the traffic between the hosts (end to end) and check the usage.
7. Check the duplex and speed settings in ASA interfaces. The mismatch setting with the remote interfaces can increase the CPU utilization.

This example shows the higher number in *input error* and *overruns* due to the speed mismatch. Use the `show interface` command in order to verify the errors:


```
<#root>
Ciscoasa#
sh int GigabitEthernet0/1
```

```
Interface GigabitEthernet0/1 "inside", is up, line protocol is up
  Hardware is i82546GB rev03, BW 1000 Mbps, DLY 10 usec
    Auto-Duplex(Full-duplex), Auto-Speed(100 Mbps)
    Input flow control is unsupported, output flow control is unsupported
    MAC address 0013.c480.b2b8, MTU 1500
    IP address 192.168.17.4, subnet mask 255.255.255.0
    311981 packets input, 20497296 bytes, 0 no buffer
    Received 311981 broadcasts, 157 runts, 0 giants
```

```
7186 input errors, 0 CRC, 0 frame, 7186 overrun
```


```
, 0 ignored, 0 abort
  0 pause input, 0 resume input
  0 L2 decode drops
  121 packets output, 7744 bytes, 0 underruns
  0 pause output, 0 resume output
  0 output errors, 0 collisions, 1 interface resets
  0 late collisions, 0 deferred
  0 input reset drops, 0 output reset drops, 0 tx hangs
  input queue (blocks free curr/low): hardware (255/249)
  output queue (blocks free curr/low): hardware (255/254)
```

In order to resolve this issue, set speed as *auto* to the corresponding interface.




Note: Cisco recommends that you enable the `ip verify reverse-path interface` command on all the interfaces. This causes packets that do not have a valid source address to be dropped and results in less CPU usage. This applies to FWSM when it faces high CPU issues.

-
8. Another reason for high CPU usage can be due to too many multicast routes. Issue the [show mroute](#) command in order to check if ASA receives too many multicast routes.
 9. Use the [show local-host](#) command in order to see if the network experiences a denial-of-service attack, which can indicate a virus attack in the network.
 10. High CPU can occur due to Cisco bug ID [CSCsq48636](#) . Refer to Cisco bug ID [CSCsq48636](#) for more information.



Note: Only registered Cisco users can access internal Cisco tools and bug information.

 **Note:** If the solution provided previously does not resolve the issue, upgrade the ASA platform based on the requirements. Refer to [Cisco Security Modules for Security Appliances](#) for more information on Adaptive Security Appliance Platform capabilities and capacities. Contact TAC ([Cisco Technical Support](#)) for further information.

High Memory Utilization

Here are some possible causes and resolutions for high memory utilization:

- **Event logging:** Event logging can consume large amounts of memory. In order to resolve this issue, install and log all events to an external server, such as a syslog server.
- **Memory Leakage:** A known issue in the security appliance software can lead to high memory


consumption. In order to resolve this issue, upgrade the security appliance software.

- **Debugging Enabled:** Debugging can consume large amounts of memory. In order to resolve this issue, disable debugging with the `undebug all` command.
- **Blocking Ports:** Blocking ports on the outside interface of a security appliance cause the security appliance to consume high amounts of memory to block the packets through the specified ports. In order to resolve this issue, block the offending traffic at the ISP end.
- **Threat-Detection:** The threat detection feature consists of different levels of statistics that are gathered for various threats and scanned threat detection, which determines when a host performs a scan. **Turn off** this feature to consume less memory.

PortFast, Channeling, and Trunking

By default, many switches, such as Cisco switches that run the Catalyst operating system (OS), are designed to be plug-and-play devices. As such, many of the default port parameters are not desirable when an ASA is plugged into the switch. For example, on a switch that runs the Catalyst OS, default channeling is set to Auto, trunking is set to Auto, and PortFast is disabled. If you connect an ASA to a switch that runs the Catalyst OS, disable channeling, disable trunking, and enable PortFast.


Channeling, also known as Fast EtherChannel or Giga EtherChannel, is used to bind two or more physical ports in a logical group in order to increase the overall throughput across the link. When a port is configured for automatic channeling, it sends out Port Aggregation Protocol (PAgP) frames as the link becomes active in order to determine if it is part of a channel. These frames can cause problems if the other device tries to automatically negotiate the speed and duplex of the link. If channeling on the port is set to Auto, it also results in an additional delay of about 3 seconds before the port starts to forward traffic after the link is up.

 **Note:** On the Catalyst XL Series Switches, channeling is not set to Auto by default. For this reason, you must disable channeling on any switch port that connects to an ASA.

Trunking, also known by the common trunking protocols Inter-Switch Link (ISL) or Dot1q, combines multiple virtual LANs (VLANs) on a single port (or link). Trunking is typically used between two switches when both switches have more than one VLAN defined on them. When a port is configured for automatic trunking, it sends out Dynamic Trunking Protocol (DTP) frames as the link comes up in order to determine if the port that it connects to wants to trunk. These DTP frames can cause problems with automatic negotiation of the link. If trunking is set to Auto on a switch port, it adds an additional delay of about 15 seconds before the port starts to forward traffic after the link is up.


PortFast, also known as Fast Start, is an option that informs the switch that a Layer 3 device is connected out of a switch port. The port does not wait the default 30 seconds (15 seconds to listen and 15 seconds to learn); instead, this action causes the switch to put the port into forwarding state immediately after the link

comes up. It is important to understand that when you enable PortFast, spanning tree is not disabled. Spanning tree is still active on that port. When you enable PortFast, the switch is informed only that there is not another switch or hub (Layer 2-only device) connected at the other end of the link. The switch bypasses the normal 30-second delay while it attempts to determine if a Layer 2 loop results if it brings up that port. After the link is brought up, it still participates in spanning tree. The port sends out bridge packet data units (BPDUs), and the switch still listens for BPDUs on that port. For these reasons, it is recommended that you enable PortFast on any switch port that connects to an ASA.

 **Note:** Catalyst OS releases 5.4 and later include the `set port host <mod>/<port>` command that allows you to use a single command to disable channeling, disable trunking, and enable PortFast.

Network Address Translation (NAT)

Each NAT or NAT Overload (PAT) session is assigned a translation slot known as an *xlate*. These xlates can persist even after you make changes to the NAT rules that affect them. This can lead to a depletion of translation slots or unexpected behavior or both by traffic that undergoes translation. This section explains how to view and clear xlates on the security appliance.

 **Caution:** A momentary interruption of the flow of all traffic through the device can occur when you globally clear xlates on the security appliance.

Sample ASA configuration for PAT that uses the outside interface IP Address:

```
object network OBJ_GENERIC_ALL
subnet 0.0.0.0 0.0.0.0

nat (inside,outside) source dynamic OBJ_GENERIC_ALL interface
```

Traffic that flows through the security appliance most likely undergoes NAT. In order to view the translations that are in use on the security appliance, issue the `show xlate` command:

```
<#root>
```

```
Ciscoasa#
```

```
show xlate
```

```
5 in use, 5 most used
```

```
Flags: D - DNS, i - dynamic, r - portmap, s - static, I - identity, T - twice
```

```
NAT from any:192.168.1.10 to any:172.16.1.1/24
```

```
flags s idle 277:05:26 timeout 0:00:00
```

Translation slots can persist after key changes are made. In order to clear current translation slots on the security appliance, issue the `clear xlate` command:

```
<#root>
```

```
Ciscoasa#
```

```
clear xlate
```

```
<#root>
```

```
Ciscoasa#
```

```
show xlate
```

```
0 in use, 1 most used
```

The `clear xlate` command clears all the current dynamic translation from the xlate table. In order to clear a particular IP translation, you can use the `clear xlate` command with the `global [ip address]` keyword.

Here is a sample ASA configuration for NAT:

```
object network inside-net
subnet 0.0.0.0 0.0.0.0
object network outside-pat-pool
range 10.10.10.10 10.10.10.100
nat (inside,outside) source dynamic inside-net outside-pat-pool
```

Observe the `show xlate` output for the translation for inside 10.2.2.2 to outside global 10.10.10.10:

```
<#root>
Ciscoasa#
show xlate
2 in use, 2 most used

Flags: D - DNS, i - dynamic, r - portmap, s - static, I - identity, T -
twice

TCP PAT from inside:10.2.2.2/1429 to any:10.10.10.10/64768 flags ri
idle 62:33:57 timeout 0:00:30

TCP PAT from inside:10.5.5.5/1429 to any:10.10.10.11/64768 flags ri
idle 62:33:57 timeout 0:00:30
```

Clear the translation for 10.10.10.10 global IP address:

```
<#root>
Ciscoasa# clear xlate global 10.10.10.10
```

In this example, the translation for inside 10.2.2.2 to outside global 10.10.10.10 is gone:

```
<#root>
Ciscoasa#
show xlate
1 in use, 2 most used

Flags: D - DNS, i - dynamic, r - portmap, s - static, I - identity, T -
twice

TCP PAT from inside:10.5.5.5/1429 to any:10.10.10.11/64768 flags ri
idle 62:33:57 timeout 0:00:30
```

Syslogs


Syslogs allow you to troubleshoot issues on the ASA. Cisco offers a free syslog server for Windows NT called ASA Firewall Syslog Server (PFSS). You can download PFSS from [Cisco Technical Support & Downloads](#).

Several other vendors, such offer syslog servers for various Windows platforms, such as Windows 2000 and Windows XP. Most UNIX and Linux machines have syslog servers installed by default.

When you set up the syslog server, configure the ASA in order to send logs to it.

For example:

```
<#root>  
  
logging on  
logging host <ip_address_of_syslog_server>  
logging trap debugging
```

 **Note:** This example configures the ASA to send Debugging (level 7) and more critical syslogs to the syslog server. Because these ASA logs are the most verbose, use them only when you troubleshoot an issue. For normal operation, configure the logging level to Warning (level 4) or Error (level 3).

If you experience an issue with slow performance, open the syslog in a text file and search for the source IP address associated with the performance issue. (If you use UNIX, you can `grep` through the syslog for the source IP address.) Check for messages that indicate the external server tried to access the internal IP address on TCP port 113 (for Identification Protocol, or Ident), but the ASA denied the packet. The message must be similar to this example:

```
%ASA-2-106001: Inbound TCP connection denied from  
10.64.10.2/35969 to 192.168.110.179/113 flags SYN
```

If you receive this message, issue the `service resetinbound` command to the ASA. The ASA does not silently drop packets; instead, this command causes the ASA to immediately reset any inbound connection that is denied by the security policy. The server does not wait for the Ident packet to time out its TCP connection; instead, it immediately receives a reset packet.

SNMP

A recommended method for the enterprise deployments is to monitor the performance of Cisco ASA with SNMP. Cisco ASA supports this with SNMP versions 1, 2c, and 3.

You can configure the Security appliance to send traps to a Network Management Server (NMS), or you can use the NMS to browse the MIBs on the security appliance. MIBs are a collection of definitions, and the security appliance maintains a database of values for each definition. For more information about this, refer to [Cisco ASA 5500 Series Configuration Guide with the CLI, 8.4 and 8.6](#).

All the supported MIBs for Cisco ASA can be found at ASA MIB Support List. From this list, these MIBs are useful when you monitor performance:

- CISCO-FIREWALL-MIB ---- Contains Objects useful for failover.
- CISCO-PROCESS-MIB ---- Contains Objects useful for CPU Utilization.
- CISCO-MEMORY-POOL-MIB ---- Contains Objects useful for Memory Objects.

Reverse DNS Lookups

If you experience slow performance with the ASA, verify that you have Domain Name System Pointer (DNS PTR) records, also known as Reverse DNS Lookup records, in the authoritative DNS server for the external addresses that the ASA uses. This includes any address in your global Network Address Translation (NAT) pool (or the ASA outside interface if you overload on the interface), any static address, and internal address (if you do not use NAT with them). Some applications, such as File Transfer Protocol (FTP) and Telnet servers, can use reverse DNS lookups in order to determine where the user comes from and if it is a valid host. If the reverse DNS lookup does not resolve, then performance is degraded as the request times out.

In order to ensure that a PTR record exists for these hosts, issue the `nslookup` command from your PC or UNIX machine; include the global IP address you use to connect to the Internet.

Example

```
<#root>  
  
% nslookup 192.168.219.25  
  
10.219.133.198.in-addr.arpa      name = www.cisco.com.
```

You must receive a response back with the DNS name of the device assigned to that IP address. If you do not receive a response, contact the person that controls your DNS in order to request the addition of PTR records for each of your global IP addresses.

Overruns on the Interface

If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. If you enable pause frames for flow control can alleviate this issue. Pause (XOFF) and XON frames are generated automatically by the NIC hardware-based on the FIFO buffer usage. A pause frame is sent when the buffer usage exceeds the high-water mark. In order to enable pause (XOFF) frames for flow control, use this command:

```
<#root>  
  
hostname(config)#  
  
interface tengigabitethernet 1/0  
  
hostname(config-if)#  
  
flowcontrol send on
```

Show Commands

Show CPU Usage

The `show cpu usage` command is used to determine the traffic load placed on the ASA CPU. During peak traffic times, network surges, or attacks, the CPU usage can spike.

The ASA has a single CPU to process a variety of tasks; for example, it processes packets and prints debug messages to the console. Each process has its own purpose, and some processes require more CPU time than other processes. Encryption is probably the most CPU-intensive process, so if your ASA passes a lot of traffic through encrypted tunnels, you must consider a faster ASA, a dedicated VPN Concentrator, such as the VPN 3000. The VAC offloads the encryption and decryption from the ASA CPU and performs it in hardware on the card. This allows the ASA to encrypt and decrypt 100 Mbps of traffic with 3DES (168-bit encryption).

Logging is another process that can consume large amounts of system resources. Because of this, it is recommended that you disable console, monitor, and buffer logging on the ASA. You can enable these processes when you troubleshoot a problem, but disable them for day-to-day operation, especially if you run out of CPU capacity. It is also suggested that syslog or Simple Network Management Protocol (SNMP) logging (logging history) must be set to level 5 (Notification) or lower. In addition, you can disable specific syslog message IDs with the `no logging message <syslog_id>` command.

Cisco Adaptive Security Device Manager (ASDM) also provides a graph on the Monitoring tab that allows you to view the CPU usage of the ASA over time. You can use this graph in order to determine the load on your ASA.

The `show cpu usage` command can be used to display CPU utilization statistics.

Example

```
<#root>
```

```
Ciscoasa#
```

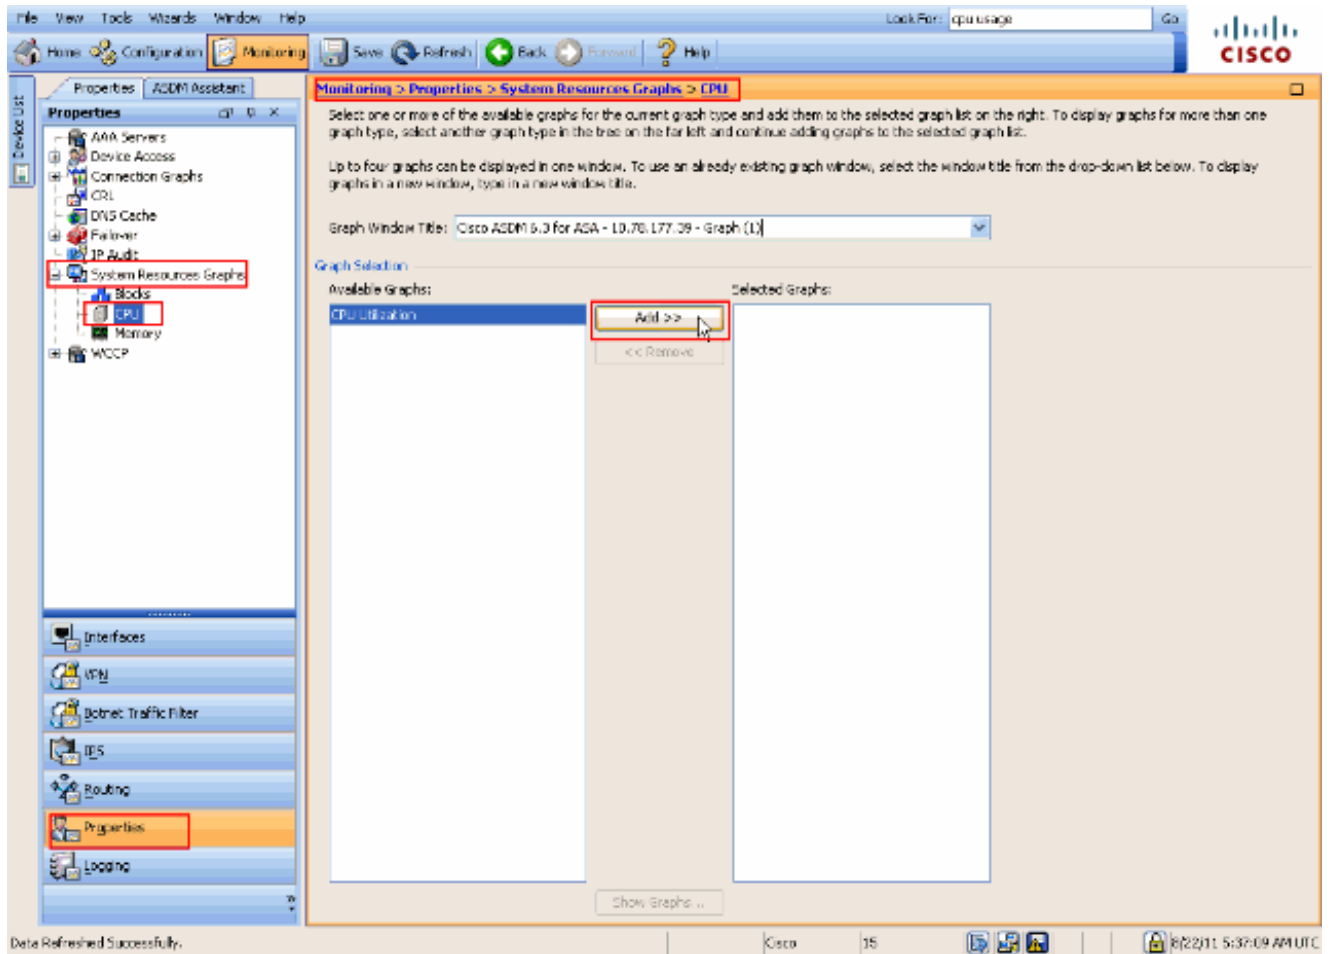
```
show cpu usage
```

```
CPU utilization for 5 seconds = 1%; 1 minute: 2%; 5 minutes: 1%
```


View CPU Usage on ASDM

Complete these steps in order to view the CPU usage on the ASDM:


1. Go to Monitoring > Properties > System Resources Graphics > CPU in ASDM and choose the **Graph Window Title**. Then, choose the required graphs from the list of **Available Graphs** and click **Add** as shown.



2. Once the required graph name is added under the **Selected Graphs** section, click **Show Graphs**.



The next image shows the **CPU Usage** graph on the ASDM. Different views of this graph are available and can be changed when the view from the View drop-down list is selected. This output can be printed or saved to the computer as required.



Description of Output

This table describes the fields in the `show cpu usage` output.

Field	Description
-------	-------------

CPU utilization for 5 seconds	CPU utilization for the last five seconds
1 minute	Average of 5 second samples of CPU utilization over the last minute
5 minutes	Average of 5 second samples of CPU utilization over the last five minutes

Show Traffic

The [show traffic](#) command shows how much traffic that passes through the ASA over a given period of time. The results are based on the time interval since the command was last issued. For accurate results, issue the [clear traffic](#) command first and then wait 1-10 minutes before you issue the `show traffic` command. You could also issue the `show traffic` command and wait 1-10 minutes before you issue the command again, but only the output from the second instance is valid.

You can use the `show traffic` command in order to determine how much traffic passes through your ASA. If you have multiple interfaces, the command can help you determine which interfaces send and receive the most data. For ASA appliances with two interfaces, the sum of the inbound and outbound traffic on the outside interface must equal the sum of the inbound and outbound traffic on the inside interface.

Example

```
<#root>
Ciscoasa#
show traffic
outside:
  received (in 124.650 secs):
    295468 packets  167218253 bytes
    2370 pkts/sec   1341502 bytes/sec
  transmitted (in 124.650 secs):
    260901 packets  120467981 bytes
    2093 pkts/sec   966449 bytes/sec
inside:
  received (in 124.650 secs):
    261478 packets  120145678 bytes
    2097 pkts/sec   963864 bytes/sec
  transmitted (in 124.650 secs):
    294649 packets  167380042 bytes
    2363 pkts/sec   1342800 bytes/sec
```

If you come close to or reach the rated throughput on one of your interfaces, you need to upgrade to a faster interface or limit the amount of traffic that goes into or out of that interface. Failure to do so can result in dropped packets. As explained in the `show interface` section, you can examine the interface counters in order to find out about throughput.

Show Perfmon

The `show perfmon` command is used to monitor the amount and types of traffic that the ASA inspects. This command is the only way to determine the number of translations (xlates) and connections (conn) per second. Connections are further broken down into TCP and User Datagram Protocol (UDP) connections. See **Description of Output** for descriptions of the output that this command generates.

Example

PERFMON STATS	Current	Average
Xlates	18/s	19/s
Connections	75/s	79/s
TCP Conns	44/s	49/s
UDP Conns	31/s	30/s
URL Access	27/s	30/s
URL Server Req	0/s	0/s
TCP Fixup	1323/s	1413/s
TCPIntercept	0/s	0/s
HTTP Fixup	923/s	935/s
FTP Fixup	4/s	2/s
AAA Authen	0/s	0/s
AAA Author	0/s	0/s
AAA Account	0/s	0/s

Description of Output

This table describes the fields in the `show perfmon` output.

Field	Description
Xlates	Translations built up per second

Connections	Connections established per second
TCP Conns	TCP connections per second
UDP Conns	UDP connections per second
URL Access	URLs (websites) accessed per second
URL Server Req	Requests sent to Websense and N2H2 per second (requires <code>filter</code> command)
TCP Fixup	Number of TCP packets that the ASA forwards per second
TCP Intercept	Number of SYN packets per second that have exceeded the embryonic limit set on a static
HTTP Fixup	Number of packets destined to port 80 per second (requires <code>fixup protocol http</code> command)
FTP Fixup	FTP commands inspected per second
AAA Authen	Authentication requests per second
AAA Author	Authorization requests per second
AAA Account	Accounting requests per second

Show Blocks

Along with the `show cpu usage` command, you can use the [show blocks](#) command in order to determine whether the ASA is overloaded.

Packet Blocks (1550 and 16384 Bytes)

When it comes into the ASA interface, a packet is placed on the input interface queue, passed up to the OS, and placed in a block. For Ethernet packets, the 1550-byte blocks are used; if the packet comes in on a 66 MHz Gigabit Ethernet card, the 16384-byte blocks are used. The ASA determines whether the packet is permitted or denied based on the Adaptive Security Algorithm (ASA) and processes the packet through to

the output queue on the outbound interface. If the ASA cannot support the traffic load, the number of available 1550-byte blocks (or 16384-byte blocks for 66 MHz GE) hovers close to 0 (as shown in the CNT column of the command output). When the CNT column hits zero, the ASA attempts to allocate more blocks, up to a maximum of 8192. If no more blocks are available, the ASA drops the packet.

Failover and Syslog Blocks (256 Bytes)

The 256-byte blocks are mainly used for stateful failover messages. The active ASA generates and sends packets to the standby ASA in order to update the translation and connection table. During periods of bursty traffic where high rates of connections are created or torn down, the number of available 256-byte blocks can drop to 0. This drop indicates that one or more connections are not updated to the standby ASA. This is generally acceptable because the next time around the stateful failover protocol catches the xlate or connection that is lost. However, if the CNT column for 256-byte blocks stays at or near 0 for extended periods of time, the ASA cannot keep up with the translation and connection tables that are synchronized because of the number of connections per second that the ASA processes. If this happens consistently, upgrade the ASA to a faster model.

Syslog messages sent out from the ASA also use the 256-byte blocks, but they are not generally released in such a quantity that causes a depletion of the 256-byte block pool. If the CNT column shows that the number of 256-byte blocks is near 0, ensure that you do not log at Debugging (level 7) to the syslog server. This is indicated by the logging trap line in the ASA configuration. It is recommended that you set logging to Notification (level 5) or lower unless you require additional information for debugging purposes.

Example

```
<#root>
```

```
Ciscoasa#
```

```
show blocks
```

SIZE	MAX	LOW	CNT
4	1600	1597	1600
80	400	399	400
256	500	495	499
1550	1444	1170	1188
16384	2048	1532	1538

Description of Output

This table describes the columns in the `show blocks` output.

Column	Description
SIZE	E Size, in bytes, of the block pool. Each size represents a particular type
MAX	Maximum number of blocks available for the specified byte block pool. The maximum number of blocks are carved out of memory at bootup. Typically, the maximum number of blocks does not change. The exception is for the 256- and 1550-byte blocks, where the adaptive security appliance can dynamically create more when needed, up to a maximum of 8192.
LOW	Low-water mark. This number indicates the lowest number of this size blocks available since the adaptive security appliance was powered up, or since the last clearing of the blocks (with the clear blocks command). A zero in the LOW column indicates a previous event where memory was full.
CNT	Current number of blocks available for that specific size block pool. A zero in the CNT column means memory is full now.

This table describes the SIZE row values in the `show blocks` output.

SIZE Value	Description
0	Used by dupb blocks.
4	Duplicates existing blocks in applications such as DNS, ISAKMP, URL filtering, uauth, TFTP, and TCP modules. Also, this sized block can be used normally by code to send packets to drivers, and so on.
80	Used in TCP intercept to generate acknowledgment packets and for failover hello messages.
256	Used for Stateful Failover updates, syslog logging, and other TCP functions. These blocks are mainly used for Stateful Failover messages. The active adaptive security appliance generates and sends packets to the standby adaptive security appliance to update the translation and connection table. In bursty traffic, where high rates of connections are created or torn down, the number of available blocks can drop to 0. This situation indicates that one or more connections were not updated to the standby adaptive security appliance. The Stateful Failover protocol catches the lost translation or connection the next time. If the CNT column for 256-byte blocks stays at or near 0 for extended periods of time, then the adaptive security appliance struggles to keep up the translation and connection tables synchronized because of the number of connections per second that the

	adaptive security appliance processes. Syslog messages sent out from the adaptive security appliance also use the 256-byte blocks, but they are generally not released in such quantity to cause a depletion of the 256-byte block pool. If the CNT column shows that the number of 256-byte blocks is near 0, ensure that you are not logging at Debugging (level 7) to the syslog server. This is indicated by the logging trap line in the adaptive security appliance configuration. We recommend that you set logging at Notification (level 5) or lower unless you require additional information for debugging purposes.
1550	Used to store Ethernet packets to process through the adaptive security appliance. When a packet enters an adaptive security appliance interface, it is placed on the input interface queue, passed up to the operating system, and placed in a block. The adaptive security appliance determines whether the packet must be permitted or denied based on the security policy and processes the packet through to the output queue on the outbound interface. If the adaptive security appliance struggles to keep up with the traffic load, the number of available blocks can hover close to 0 (as shown in the CNT column of the command output). When the CNT column is zero, the adaptive security appliance attempts to allocate more blocks, up to a maximum of 8192. If no more blocks are available, the adaptive security appliance drops the packet.
16384	Only used for the 64-bit, 66-MHz Gigabit Ethernet cards (i82543). See the description for 1550 for more information about Ethernet packets.
2048	Control or guided frames used for control updates.

Show Memory

The [show memory](#) command displays the total physical memory (or RAM) for the ASA, along with the number of bytes currently available. In order to use this information, you must first understand how the ASA uses memory. When the ASA boots, it copies the OS from Flash into RAM and runs the OS from RAM (just like routers). Next, the ASA copies the startup configuration from Flash and places it into RAM. Finally, the ASA allocates RAM in order to create the block pools discussed in the [show blocks](#) section. Once this allocation is complete, the ASA needs additional RAM only if the configuration increases in size. In addition, the ASA stores the translation and connection entries in RAM.

During normal operation, the free memory on the ASA must change very little, if at all. Typically, the only time you must run low on memory is if you are under attack and hundreds of thousands of connections go through the ASA. In order to check the connections, issue the `show conn count` command, which displays the current and maximum number of connections through the ASA. If the ASA runs out of memory, it eventually crashes. Prior to the crash, you can notice memory allocation failure messages in the syslog (%ASA-3-211001).

If you run out of memory because you are under attack, contact the [Cisco Technical Support](#) team.


Example



```
<#root>
Ciscoasa#
show memory

Free memory:      845044716 bytes (79%)
Used memory:      228697108 bytes (21%)
-----
Total memory:     1073741824 bytes (100%)
```

Show Xlate

The [show xlate count](#) command displays the current and maximum number of translations through the ASA. A translation is a mapping of an internal address to an external address and can be a one-to-one mapping, such as Network Address Translation (NAT), or a many-to-one mapping, such as Port Address Translation (PAT). This command is a subset of the `show xlate` command, which outputs each translation through the ASA. Command output shows translations "in use," which refers to the number of active translations in the ASA when the command is issued; "most used" refers to the maximum translations that have ever been seen on the ASA since it was powered on.

 **Note:** A single host can have multiple connections to various destinations, but only one translation. If the xlate count is much larger than the number of hosts on your internal network, it is possible that one of your internal hosts has been compromised. If your internal host has been compromised, it spoofs the source address and sends packets out the ASA.

 **Note:** When the vpnclient configuration is enabled and the inside host sends out DNS requests, the `show xlate` command can list multiple xlates for a static translation.

Example

```
<#root>
Ciscoasa#
```

```
show xlate count
```

```
84 in use, 218 most used
```

```
<#root>
```

```
Ciscoasa(config)#
```

```
show xlate
```

```
3 in use, 3 most used
```

```
Flags: D - DNS, d - dump, I - identity, i - inside, n - no random,  
       o - outside, r - portmap, s - static
```

```
TCP PAT from inside:10.1.1.15/1026 to outside:192.168.49.1/1024 flags ri  
idle 62:33:57 timeout 0:00:30
```

```
UDP PAT from 10.1.1.15/1028 to outside:192.168.49.1/1024 flags ri  
idle 62:33:57 timeout 0:00:30
```

```
ICMP PAT from inside:10.1.1.15/21505 to outside:192.168.49.1/0 flags ri  
idle 62:33:57 timeout 0:00:30
```

The first entry is a TCP Port Address Translation for host-port (10.1.1.15, 1026) on the inside network to host-port (192.168.49.1, 1024) on the outside network. The "r" flag denotes the translation is a Port Address Translation. The "i" flags denotes that the translation applies to the inside address-port.

The second entry is a UDP Port Address Translation for host-port (10.1.1.15, 1028) on the inside network to host-port (192.168.49.1, 1024) on the outside network. The "r" flag denotes the translation is a Port Address Translation. The "i" flags denotes that the translation applies to the inside address-port.

The third entry is an ICMP Port Address Translation for host-ICMP-id (10.1.1.15, 21505) on the inside network to host-ICMP-id (192.168.49.1, 0) on the outside network. The "r" flag denotes the translation is a Port Address Translation. The "i" flags denotes that the translation applies to the inside address-ICMP-id.

The inside address fields appear as source addresses on packets that traverse from the more secure interface to the less secure interface. Conversely, they appear as destination addresses on packets that traverse from the less secure interface to the more secure interface.

Show Conn Count

The [show conn count](#) command shows the current and maximum number of connections through the ASA. A connection is a mapping of Layer 4 information from an internal address to an external address. Connections are built up when the ASA receives a SYN packet for TCP sessions or when the first packet in a UDP session arrives. Connections are torn down when the ASA receives the final ACK packet, which occurs when the TCP session handshake closes or when the timeout expires in the UDP session.

Extremely high connection counts (50-100 times normal) can indicate that you are under attack. Issue the `show memory` command in order to ensure that the high connection count does not cause the ASA to run out of memory. If you are under attack, you can limit the maximum number of connections per static entry and also limit the maximum number of embryonic connections. This action protects your internal servers, so they do not become overwhelmed. Refer to [Cisco ASA 5500 Series Configuration Guide with the CLI, 8.4 and 8.6](#) for more information.

Example

```
<#root>
Ciscoasa#
show conn count
2289 in use, 44729 most used
```

Show Interface

The [show interface](#) command can help determine duplex mismatch problems and cable issues. It can also provide further insight as to whether or not the interface is overrun. If the ASA runs out of CPU capacity, the number of 1550-byte blocks hovers close to 0. (Look at the 16384-byte blocks on the 66 MHz Gig cards.) Another indicator is the increase of "no buffers" on the interface. The no buffers message indicates that the interface is unable to send the packet to the ASA OS because there is no available block for the packet, and the packet is dropped. If an increase in no buffer levels occurs regularly, issue the `show proc cpu` command in order to check the CPU usage on the ASA. If the CPU usage is high because of a heavy traffic load, upgrade to a more powerful ASA that can handle the load.

When a packet first enters an interface, it is placed in the input hardware queue. If the input hardware queue is full, the packet is placed in the input software queue. The packet is passed from its input queue and placed in a 1550-byte block (or in a 16384-byte block on 66 MHz Gigabit Ethernet interfaces). The ASA then

determines the output interface for the packet and places the packet in the appropriate hardware queue. If the hardware queue is full, the packet is placed in the output software queue. If the maximum blocks in either of the software queues are large, then the interface is overrun. For example, if 200 Mbps come into the ASA and all go out a single 100 Mbps interface, the output software queue indicates high numbers on the outbound interface, which indicates that the interface cannot handle the traffic volume. If you experience this situation, upgrade to a faster interface.

Example

```
<#root>
Ciscoasa#
show interface

Interface GigabitEthernet0/1 "inside", is up, line protocol is up
  Hardware is i82546GB rev03, BW 1000 Mbps, DLY 10 usec
    Auto-Duplex(Full-duplex), Auto-Speed(100 Mbps)
    Input flow control is unsupported, output flow control is unsupported
    MAC address 0013.c480.b2b8, MTU 1500
    IP address 192.168.17.4, subnet mask 255.255.255.0
    311981 packets input, 20497296 bytes, 0 no buffer
    Received 311981 broadcasts, 157 runts, 0 giants
    379 input errors, 107 CRC, 273 frame, 0 overrun, 0 ignored, 0 abort
    0 pause input, 0 resume input
    0 L2 decode drops
    121 packets output, 7744 bytes, 0 underruns
    0 pause output, 0 resume output
    0 output errors, 0 collisions, 1 interface resets
    0 late collisions, 0 deferred
    0 input reset drops, 0 output reset drops, 0 tx hangs
    input queue (blocks free curr/low): hardware (255/249)
    output queue (blocks free curr/low): hardware (255/254)
```


You must also check the interface for errors. If you receive runts, input errors, CRCs, or frame errors, it is likely that you have a duplex mismatch. The cable can be faulty as well. See [Speed and Duplex Settings](#) for more information on duplex issues. Remember that each error counter represents the number of packets that are dropped because of that particular error. If you see a specific counter that increments regularly, the performance on your ASA most likely suffers, and you must find the root cause of the problem.

While you examine the interface counters, note that if the interface is set to full-duplex, you must not experience any collisions, late collisions, or deferred packets. Conversely, if the interface is set to half-duplex, you must receive collisions, some late collisions, and possibly some deferred packets. The total number of collisions, late collisions, and deferred packets must not exceed 10% of the sum of the input and output packet counters. If your collisions exceed 10% of your total traffic, then the link is overutilized, and you must upgrade to full-duplex or to a faster speed (10 Mbps to 100 Mbps). Remember that collisions of 10% mean that the ASA drops 10% of the packets that go through that interface; each of these packets must be retransmitted.

Refer to the `interface` command in [Cisco ASA 5500 Series Adaptive Security Appliances Command References](#) for detailed information on the interface counters.

Show Processes

The `show processes` command on the ASA displays all the active processes that run on the ASA at the time the command is executed. This information is useful in order to determine which processes receive too much CPU time and which processes do not receive any CPU time. In order to get this information, issue the `show processes` command twice; wait about 1 minute between each instance. For the process in question, subtract the Runtime value displayed in the second output from the Runtime value displayed in the first output. This result shows you how much CPU time (in milliseconds) the process received in that interval of time. Note that some processes are scheduled to run at particular intervals, and some processes only run when they have information to process. The `577poll` process most likely has the largest Runtime value of all your processes. This is normal because the `577poll` process polls the Ethernet interfaces in order to see if they have any data that needs to be processed.

 **Note:** An examination of each ASA process is out of the scope of this document but is mentioned briefly for completeness. Refer to [ASA 8.3 and Later: Monitor and Troubleshoot Performance Issues](#) for more information about the ASA processes.

Command Summary

In summary, use the `show cpu usage` command in order to identify the load that the ASA is under. Remember that the output is a running average; the ASA can have higher spikes of CPU usage that are masked by the running average. Once the ASA reaches 80% CPU usage, the latency through the ASA slowly increases to about 90% CPU. When CPU usage is more than 90%, the ASA starts to drop packets.

If the CPU usage is high, use the `show processes` command in order to identify the processes that use the most CPU time. Use this information in order to reduce some of the time that is consumed by the intensive processes (such as logging).

If the CPU does not run hot, but you believe packets are still dropped, use the `show interface` command in order to check the ASA interface for no buffers and collisions, possibly caused by a duplex mismatch. If the no buffer count increments, but the CPU usage is not low, the interface cannot support the traffic that flows through it.

If the buffers are fine, check the blocks. If the current CNT column in the `show blocks` output is close to 0 on the 1550-byte blocks (16384-byte blocks for 66 MHz Gig cards), the ASA most likely drops Ethernet packets because it is too busy. In this instance, the CPU spikes high.

If you experience trouble when you make new connections through the ASA, use the `show conn count` command in order to check the current count of connections through the ASA.

If the current count is high, check the `show memory` output in order to ensure that the ASA does not run out of memory. If memory is low, investigate the source of the connections with the `show conn` or `show local-host` command in order to verify that your network has not experienced a denial-of-service attack.

You can use other commands in order to measure the amount of traffic that passes through the ASA. The `show traffic` command displays the aggregate packets and bytes per interface, and the `show perfmon` breaks the traffic down into different types that the ASA inspects.

Related Information

- [Cisco ASA 5500-X Series Firewalls](#)
- [Cisco Technical Support & Downloads](#)